



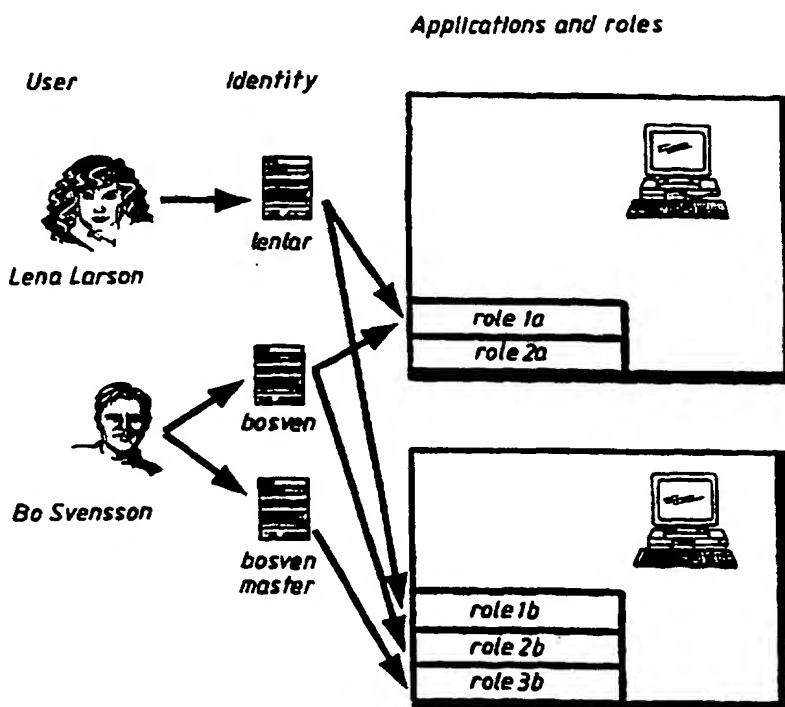
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 1/00, 12/14		A2	(11) International Publication Number: WO 96/17286
		(43) International Publication Date: 6 June 1996 (06.06.96)	
(21) International Application Number: PCT/SE95/01394		(81) Designated States: US, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 22 November 1995 (22.11.95)		Published <i>Without international search report and to be republished upon receipt of that report.</i>	
(30) Priority Data: 9404157-1 29 November 1994 (29.11.94) SE			
(71) Applicant (for all designated States except US): TELIA AB [SE/SE]; S-123 86 Farsta (SE).			
(72) Inventor; and (75) Inventor/Applicant (for US only): BIGGE, Peter [SE/SE]; Styrmansgatan 17B, S-114 54 Stockholm (SE).			
(74) Agent: KARLSSON, Berne; Telia Research AB, Rudsjötterrassen 2, S-136 80 Haninge (SE).			

(54) Title: A METHOD FOR CONTROLLING ACCESS TO A DATA BASE, A DATA BASE AND A COMPUTER NETWORK USING THE SAME

(57) Abstract

The present invention solves the problem of access control in a large data base with a large number of users by defining access rights separately for each application run on the data base. A plurality of data base roles are defined, in terms of the basic data base procedures CRUD, (create, read, update and delete), which a user is entitled to use in respect of the various value sets stored in the data base. A given user is given access to a set of roles on an application-by-application basis. When a user logs-in to a particular data base application, to which he has access, a procedure is run by the file server on which the data base is stored. The file server generates the access rights for that user in respect of that application. The user rights are then stored in the file server, preferably on a cache memory, for the time in which the application in question is running on the file server. Preferably, when the user logs-out of that application the access rights are deleted from store.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

A METHOD FOR CONTROLLING ACCESS TO A DATA BASE,
A DATA BASE AND A COMPUTER NETWORK USING THE SAME

5 The present invention relates to an improved procedure for defining user access rights to a data base, data bases using said procedure and computer networks on which such databases may be operated.

Administration of large data bases with large numbers of users has a number of problems associated with operation of the data base. Consider a data base which:

- 10 - uses a single file server as a platform;
 - has a large number of users, of the order of 20,000, half of which are simultaneously logged on to the system;
 - has a single level of administration security;
 - 15 - operates various data base procedures each of which may include some functions denied to certain users; and
 - contains data on various subjects, having diverse security ratings, some of which may only be available to a subset of data base users.
- 20

25 Users of such a data base will normally have a need to use several data base applications in parallel. Users may have restricted access rights as between different data base subjects, or even within a subject. For example, a given user may be allowed access to telephone number lists for the Stockholm area, but be denied access to data on financial ledgers. Certain

- 2 -

5 telephone numbers may be regarded as secret and access
to such numbers may be denied to all except a few users.
If there is only a single level of administration
available, the system administrator role is the only
means by which new identities can be created.
Unfortunately, this role is able to access all
information on the system, erase portions of the data,
or even the entire data base, or even switch off the
file server on which the data base is stored. The
10 problem of controlling access to such a data base,
without imposing an excessive computational load on a
file server and, thereby seriously slowing down its
speed of operation, is by no means easy to solve.

15 In order to achieve this, the access rights of each
of some 20,000 users must be broken down:

- by class of subject matter;
- within subject matter classes, by secret data;
and
- by data base operation.

20 During the performance of a given data base task
for a given user, different access rights may come into
play.

25 If effective data base security is to be preserved,
it is vital that all data base access rules be held on
the file server. If such rules were held on client
controlled equipment they could be subverted by a
dishonest client.

30 It has been proposed that management of client
identities and access rights be controlled in front of
the file server. Such front-end based systems establish

- 3 -

a number of parallel connections in relation to general identity. This creates overhead problems, performance problems and log-in problems for the file server.

5 Where a data base can be accessed by a small number of users only, from ten to a few hundred, a number of small fixed groups of users may be established and identities associated with these groups. This solution is not viable where large numbers of users are involved.

10 The present invention solves the problem of access control in a large data base with a large number of users by defining access rights separately for each application run on the data base. A plurality of data base roles are defined, in terms of the basic data base procedures CRUD, (create, read, update and delete),
15 which a user is entitled to use in respect of the various value sets stored in the data base. A given user is given access to a set of roles on an application-by-application basis. When a user logs-in to a particular data base application, to which he has
20 access, a procedure is run by the file server on which the data base is stored. The file server generates the access rights for that user in respect of that application. The user rights are then stored in the file server, preferably on a cache memory, for the time
25 in which the application in question is running on the file server. Preferably, when the user logs-out of that application the access rights are deleted from store.

30 It should be noted that with a data base, of the type to which this invention relates, which is subject to heavy user demands, it may be advantageous to inhibit command mode operation in real time. None-the-less, the present invention may still be applicable to data base systems which can be operated in command mode.

According to a first aspect of the present invention, there is provided a computer network, having a central processing means including a file server on which a data base is stored, a plurality of terminals adapted to communicate with said central processor, said computer network being so arranged that a large plurality of users has access to said data base, and which data base contains value sets having a plurality of security ratings in respect of which a plurality of data base applications can be operated, characterised in that, there are provided access means for controlling access rights to the data base, for each user, said access means including first storage means on which are stored a plurality of data base roles, said access means being arranged to define access rights for each user by combining a selected number of roles with each data base application to which a user is permitted access, generator means for automatically generating a set of access rights for a given user when said user logs-in to a given application, and second storage means for storing details of said access rights on a file server memory at least while a data base application to which said access rights relate is running on said file server.

Preferably deletion means are provided for deleting said access rights from said second storage means when said user logs-out from said data base application.

Preferably said second storage means is a cache memory.

Preferably the number of roles is substantially less than the number of users.

According to a second aspect of the present invention there is provided a data base, stored on a

- 5 -

file server, to which a large number of users has access, containing value sets having a plurality of security ratings in respect of which a plurality of data base applications can be operated, characterised in that, access rights to the data base, for each user, are determined by establishing a plurality of data base roles, defining access rights for each user by combining a selected number of roles with each data base application to which a user is permitted access, said file server automatically generating a set of access rights for a given user when said user logs-in to a given application, storing details of said access rights on a file server memory by at least while a data base application to which said access rights relate is running on said file server.

Preferably said access rights are deleted from said file server memory when said user logs-out from said data base application.

Said file server memory may be a cache memory.

Preferably certain roles enable a user assigned those roles to perform at least some system administration functions.

Preferably said roles are defined in terms of system procedures which can be performed in relation to the different value sets held in said data base.

According to a third aspect of the present invention, in a computer network, having a central processing means including a file server on which a data base is stored, a plurality of terminals adapted to communicate with said central processor, said computer network being so arranged that a large plurality of users has access to said data base, and said data base

- 6 -

5 containing value sets having a plurality of security ratings in respect of which a plurality of data base applications can be operated, there is provided a method of setting access rights for each of said plurality of users characterised by performance of the following steps:

- establishing a plurality of data base roles;
- 10 - defining access rights for each user by combining a selected number of roles with each data base application to which a user is permitted access;
- automatically generating, with said file server, a set of access rights for a given user when said user logs-in to a given application; and
- 15 - storing details of said access rights on a file server memory while a data base application to which said access rights relate is running on said file server.

20 Preferably said access rights are automatically deleted from said file server memory when said user logs-out from said data base application.

25 Preferably the access rights for a given user, in respect of a given data base application, are determined by running a data base procedure.

Embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings in which:

30 Figure 1 illustrates the relationship established

- 7 -

between user identities, roles and applications in a data base according to the present invention.

Figure 2 illustrates how a defined role determines the access rights of a given user to a given application on a database according to the present invention.

Figure 3 shows a signal chart for an embodiment of the present invention operating on a SQL file server.

Figure 4 illustrates, in diagrammatic form, an example of a platform on which a data base according to the present invention may be run.

An example of a computer network platform on which the present invention can be realised is illustrated in Figure 4. The operation of such a platform will be immediately apparent, to those skilled in the art, from a brief examination of Figure 4. However, it should be noted that a plurality of terminals using a variety of operating systems, collectively referred to as the application system TS, can be connected to a central file server, the SQL server, by means of Sybase communications software which is herein referred to as Open Client. A variety of systems, referred to as Open Servers, may be used for preparing data base procedures and applications. The Open Servers again communicate with the central server by means of Sybase communications software. The central SQL server carries all the data which comprises the data base, and operates on SQL procedure language. That portion of the computer network which includes the central SQL server and the Open Servers is collectively referred to as the information system, IS. The open servers and applications system may use a variety of operating

- 8 -

systems and software, including UNISYS/TIP, VAX/VMS, PC, MAC, UNIX, and IBM/CICS.

5 In the operation of a data base, different roles can be identified. In a simple database, there exists two roles, that of system administrator and that of a user who wishes to access data. However, in a very large database, of the type that may be needed to support the operations of a telecommunications network operator, there may be a need for many users to perform data base applications which involve, in addition to reading data, amending data on the data base, deleting data on the database, or creating new data. The operations of creating, reading, updating and deleting data are frequently referred to by their initial letters as CRUD.

10 The relationship between a user, a user identity, a role and an application are illustrated in schematic form in Figures 1 and 2. Consider two individual users, of the very large number of total users of the database, perhaps as many as 10,000, called Lena Larson and Bo Svensson. Lena's system identity is "lenlar", and it is unique to her, Bo's computer identity is "bosven" and it is again unique to him. In addition, Bo has a second computer identity, "bosvem-master", which he uses only when accessing certain secret data for which he has special authority to access. For each identity there exists a secret password known only to the user. By entering their computer names and passwords, Bo and Lena get very limited access to the SQL file server.

25 To obtain the ability to run an application on the SQL file server, a further log-in procedure is performed automatically by the SQL file server. This will be described later. Of the many applications that can be run on the SQL file server, two are referred to, in the

- 9 -

drawings, as application A and application B. For each of these applications a number of different roles can be defined, e.g. roles 1, 2 and 3. Performance of a given application may involve a number of separate data base procedures, accessed, for example, via a menu.

Consider role 3B, that is to say role 3 performed in respect of application B. It is assumed that this application involves three operations OPa, OPb, and OPc. These three operations each use a different set of stored procedures SP1 SP4. Each of these procedures involves a number of basic data base functions, e.g. SP1 requires the create, read, update and delete functions, while SP3 involves only the delete function. By applying role 3 to application B which operates on data relating to telephone numbers, a series of actions and access rights are defined. For example, role 3B, which is only available to Bo when using the computer identity "bosvem-master", enables secret telephone numbers to be read, but not created, updated, or deleted. By way of contrast, role 3B enables telephone numbers for Northern Region to be created, read and updated, but not deleted.

In use, users, such as Lena and Bo have to be registered on the open file server. This process can only be performed by the system administrator role. In order to register a user, the system administrator enters his/her personal information such as name, organisation, and telephone number, and creates a basic, or computer, identity for the user e.g. lenlar. In addition, the system administrator assigns each user a password, which the user will have to change to a secret password on first logging in to the system. Each computer identity is assigned a run time which indicates the time of day when that user is permitted to use the system.

- 10 -

To set up access rights for a user, the system administrator has to associate a number of roles with each data base application that the user is allowed to access. As explained above, the roles define precisely what access rights a particularly user has when running a given data base application. A given individual may be assigned more than one computer identity each of which has different roles associated with it. This is advantageous when a single user may need to operate on data having different security classifications and assists in the construction of audit trails.

Access to the data base held on the information system, see Figure 4, is achieved by means of stored procedures. These procedures can create, read, update and/or delete data on the information system. In a given application, a number of operations are defined. An operation is the basic functionality seen by a user. Operations are usually selected from a menu of alternatives, or selected by operation of a specially assigned terminal key. Running a given operation either initiates a dialogue with the IS, or achieves the performance of a task in the IS. Each operation uses at least one stored procedure. These procedures are stored on the open file server. When a role is defined for an application it determines what a user who is assigned that role can do when running that application. If, for example, a user is permitted to open a dialogue with the IS, but only read the contents of stored data, this is indicated by "R", see Figure 2. Thus, for operation OPa, a user assigned role 3B is permitted to run procedures which only entail reading data, i.e. SP2. On the other hand, the user is not permitted to run procedure SP1, because this procedure requires the ability to create, read, update and delete, and role 3B is denied the right to create, update and delete within this application.

- 11 -

5 A value set is a set of criteria by which data objects are grouped. For example, all telephone subscribers in the Stockholm region represents a value set. When a role is defined, the definition includes the value sets which a user assigned that role is permitted to access, and what basic data base functions that role is permitted in respect of that value set.

10 In real time operation, the open server must deal with many users and the information has to be protected from unauthorised access yet be available to those who are authorised to use it. The definition of different roles enables users to be given varied access rights, depending on their particular needs. Equally, it is possible to arrange for a hierarchy of system administrators, by defining different system administration roles. In this way it is possible to have a large number of system administrators with only a few having unlimited access rights to the system, the majority having limited rights.

20 The totality of system users can be divided into client groups, each client representing a department, or external organisation. The use of varied administration roles can, with advantage, be implemented in a client-user hierarchy, by creating a system administrator for each client grouping with access rights limited to the needs of the client. Any changes needed to that clients overall access rights can then be achieved by a variation of the client-administrator's access rights, which must be made by one of the few central administrators. Subsequent changes to the individual user rights of users belonging to that client grouping can be implemented by the client-administrator.

The majority of system administrators thus have limited access rights, and this can be shown by the use

- 12 -

of menus, or screen buttons, in which certain menu selections, or buttons, are faded, corresponding to menu operations denied a particular system administrator. The principle reason for showing all menu options, whether available to a particular system administrator, or not, is that it enables a common set of instruction manuals and other literature to be prepared for the system. Such documentation must, of course describe maximum functionality. Of course, it is not only system operations that may be denied to a "junior" system administrator, but also certain value sets. In other words the information that a client based administrator may access can be limited to those value sets which the client is entitled to access.

In the above description of the relationship of identities, users, roles, clients, applications, value sets and system administrators, the emphasis has been on the concepts and functionality of the system, rather than on the physical means by which that functionality is created. In general, in a system of the type described, functionality is created by means of software, i.e computer programs, data base procedures, or the like. It is, of course, possible for functionality to be created by hardware elements, or means, although except in a few cases it is not cost effective to do this. The nature of the inventions described in this specification is such that the inventive concepts reside in the methods and functionality that underlies the software. From the point of view of a practical realisation of the invention it matters little whether a "means" for performing a particular function is a logical entity, created by a computer program, or a physical entity, created by a manufacturing process.

The invention(s) herein described relate to

- 13 -

improvements in real products and have a real and significant effect on product performance.

5 In general, the "means" described in this specification can readily be realised by a skilled computer programmer, familiar with modern programming techniques, from the functional/conceptual description herein provided.

10 The distinct and practical "real world" advantages of the user-role-application concept described above, together with the concept of a hierarchy of system administrators, are that they enable massive data bases, with large numbers of interactive system users, to be managed in a manner which preserves system integrity and data security. It enables the system to operate at
15 substantially greater speeds and with a substantially reduced processing overhead than is the case with other known systems.

20 As explained above, an individual user is assigned access rights on the basis of the roles he is permitted to perform in different applications. It would, in theory, be possible to provide each user with a plurality of computer identities and passwords, one for each application he is entitled to use. This is, of course, impractical, especially when considered in
25 relation to control access to different value sets. The present invention overcomes this problem by using a double log-in procedure.

30 Consider a particular user called Anders Svensson, he has a computer identity, known to himself, called "andsve" and a secret password which only he knows "pendel01". There is only a very limited amount of information on the file server that can be accessed, or changed, by the computer identity "andsve" using a

- 14 -

5 general program such as Excel Q+E. However, the double
log-in procedure give access to a much wider range of
data, on an application-by-application basis. The
double log-in procedure automatically generates, within
the file server, a plurality of computer identities
linked to "andsve", together with associated passwords,
one for each application that "andsve" is permitted to
access. These user/application identities are not
visible to Anders Svensson, and he is not permitted
10 access to them. The only log-in identity and password
known to Anders Svensson is "andsve" and "pendel01".
The log-in procedure can be written in procedure
language as:

```
15 Login(andsve,pendel01)
   Exec Fcs_get_password_for_application-identity
   logout(andsve)
   login ( a n d s v e # A P P L I C A T I O N ) ,
   password_for_application-identity
   ....Normal function calls from client
20 logout(andsve#APPLICATION)
```

The application specific identity is composed of
the identity known to the user, "andsve", and a code of
eight characters for the client separated with a "#",
(other separators could of course be used).
25 Fcs_get_password_for_application-identity creates an
application specific identity with a randomly generated
password. This identity is given rights to perform all
necessary functions in dependence on the rights of the
user in this particular application.

30 A full signal chart for the operation of the log-in
is shown in Figure 3. In the example, illustrated by
Figure 3, the basic, or computer identity "andsve",
password "pendl01", and the application code CUSTOMER
are used. Note that all error handling has been omitted
35 and that only parts of the messages of the log-in and

- 15 -

log-out sequences are indicated. Descriptions in *italic* mean that the system administrator role is used, otherwise "andsve" and "andsve#CUSTOMER" are used. Brief notes on some of the signals shown in the signal chart are set out below. However, for most steps Figure 3 is self explanatory.

Instructions *i*, *i*+1, see Figure 3, are executed once at the start of the open server.

3:- Application_LasApplid(?) initiates the creation of an application specific identity. Note that all parameters are not described in the signal chart.

4:- Check that the user is defined for the indicated application, that current time is within the users allocated run time, and that there are no restrictions on the identity, or application, etc.

6, 7:- There is a general code for access control in the open server. A check is performed back into the SQL server which verifies that the user is allowed to call. In order to reduce the access time (from about 0.3 seconds to < lms) the information is then saved in a cache memory in the open server. This cache is cleared out only when a change in the access right is introduced.

8:- The call is made with the system administrator identity, to which the open server has access. There is no possibility of the user discovering this identity and thereby gaining a higher level access than the one for which he is authorised.

9:- The identity is created and information needed to generate views of the value sets are built together with other information in the SQL server.

11:- Again "andsve" is the identity for the signal.

12:- This is the result of the function called at step 3. The result contains the randomly generated password for andsve#CUSTOMER. It is constituted by 28 characters [32255], which in addition have been DES encrypted with the password for "andsve" as key.

The remaining steps in the signal chart, illustrated in Figure 3, show the identity "andsve" being logged-out, the identity "andsve#CUSTOMER" being logged-in, the execution of the called application, "CUSTOMER", and the log-out of the identity "andsve#CUSTOMER" on completion of the application "CUSTOMER".

It should be noted that the identity "andsve#CUSTOMER" contains information on both the user and the client group to which he belongs. This form of identity is a useful aid in the derivation of audit trails, statistics and debiting etc.. In comparison, prior art systems using front-end log-in control only permit breakdowns using a single common identity, management information can be derived on a user-by-application basis. For example it is possible to enable a user to obtain the answer time of the SQL server for the application he is running.

CLAIMS

1. A computer network, having a central processing means including a file server on which a data base is stored, a plurality of terminals adapted to communicate
5 with said central processor, said computer network being so arranged that a large plurality of users has access to said data base, which data base contains value sets having a plurality of security ratings in respect of which a plurality of data base applications can be
10 operated, characterised in that, there are provided access means for controlling access rights to the data base, for each user, said access means including first storage means on which are stored a plurality of data base roles, said access means being arranged to define
15 access rights for each user by combining a selected number of roles with each data base application to which a user is permitted access, generator means for automatically generating a set of access rights for a given user when said user logs-in to a given
20 application, and second storage means for storing details of said access rights on a file server memory at least while a data base application to which said access rights relate is running on said file server.

2. A computer network as claimed in claim 1,
25 characterised in that deletion means are provided for deleting said access rights from said second storage means when said user logs-out from said data base application.

3. A computer network as claimed in claim 1, or 2,
30 characterised in that said second storage means is a cache memory.

4. A computer network as claimed in any previous

- 18 -

claim, characterised in that more than 500 users have access to said data base.

5 5. A computer network as claimed in claim 4, characterised in that more than 5,000 users have access to said data base.

6. A computer network as claimed in either claim 4, or 5, characterised in that the number of roles is substantially less than the number of users.

10 7. A computer network as claimed in any previous claim, characterised in that certain roles enable a user, assigned those roles, to perform at least some system administration functions.

15 8. A computer network as claimed in any previous claim, characterised in that said roles are defined in terms of system procedures which can be performed in relation to the different value sets held in said data base.

20 9. A computer network as claimed in claim 8, characterised in that said system procedures are, for the purposes of role definition, classified by four basic data base operations, namely, create, read, update and delete, as applied to given value sets.

25 10. A computer network as claimed in any previous claim characterised in that said access means operates by running a data base procedure which determines the access rights for a given user, in respect of a given data base application.

30 11. A data base, stored on a file server, to which a large number of users has access, containing value sets having a plurality of security ratings in respect of

- 19 -

5 which a plurality of data base applications can be
operated, characterised in that, access rights to the
data base, for each user, are determined by establishing
a plurality of data base roles, defining access rights
for each user by combining a selected number of roles
with each data base application to which a user is
permitted access, said file server automatically
generating a set of access rights for a given user when
said user logs-in to a given application, storing
10 details of said access rights on a file server memory at
least while a data base application to which said access
rights relate is running on said file server.

12. A data base as claimed in claim 11, characterised
15 in that said access rights are deleted from said file
server memory when said user logs-out from said data
base application.

13. A data base as claimed in claim 11, or 12,
characterised in that said file server memory is a cache
memory.

20 14. A data base as claimed in any of claims 11 to 13,
characterised in that more than 500 users have access to
said data base.

25 15. A data base as claimed in claim 14, characterised
in that more than 5,000 users have access to said data
base.

16. A data base as claimed in either claim 14, or 15,
characterised in that the number of roles is
substantially less than the number of users.

30 17. A data base as claimed in any of claims 11 to 16,
characterised in that certain roles enable a user
assigned those roles to perform at least some system

- 20 -

administration functions.

18. A data base as claimed in any of claims 11 to 17, characterised in that said roles are defined in terms of system procedures which can be performed in relation to the different value sets held in said data base.

19. A data base as claimed in claim 18, characterised in that said system procedures are, for the purposes of role definition, classified by four basic data base operations, namely create, read, update and delete as applied to given value sets.

20. A data base as claimed in any of claims 11 to 19, characterised in that the access rights for a given user, in respect of a given data base application are determined by running a data base procedure.

21. In a computer network, having a central processing means including a file server on which a data base is stored, a plurality of terminals adapted to communicate with said central processor, said computer network being so arranged that a large plurality of users has access to said data base, and said data base containing value sets having a plurality of security ratings in respect of which a plurality of data base applications can be operated, a method of setting access rights for each of said plurality of users characterised by performance of the following steps:

- establishing a plurality of data base roles;
- defining access rights for each user by combining a selected number of roles with each data base application to which a user is permitted access;

- 21 -

- automatically generating, with said file server, a set of access rights for a given user when said user logs-in to a given application; and
- 5 - storing details of said access rights on a file server memory while a data base application to which said access rights relate is running on said file server.

10 22. A method as claimed in claim 21, characterised by the step of automatically deleting said access rights from said file server memory when said user logs-out from said data base application.

15 23. A method as claimed in claim 21, or 22, characterised in that said file server memory is a cache memory.

24. A method as claimed in any of claims 21 to 23, characterised in that more than 500 users have access to said data base.

20 25. A method as claimed in claim 24, characterised in that more than 5,000 users have access to said data base.

26. A method as claimed in either claim 24, or 25, characterised by defining substantially fewer roles than the number of users.

25 27. A method as claimed in any of claims 21 to 26, characterised by defining some roles that enable a user assigned those roles to perform at least some system administration functions.

28. A method as claimed in any of claims 21 to 27,

- 22 -

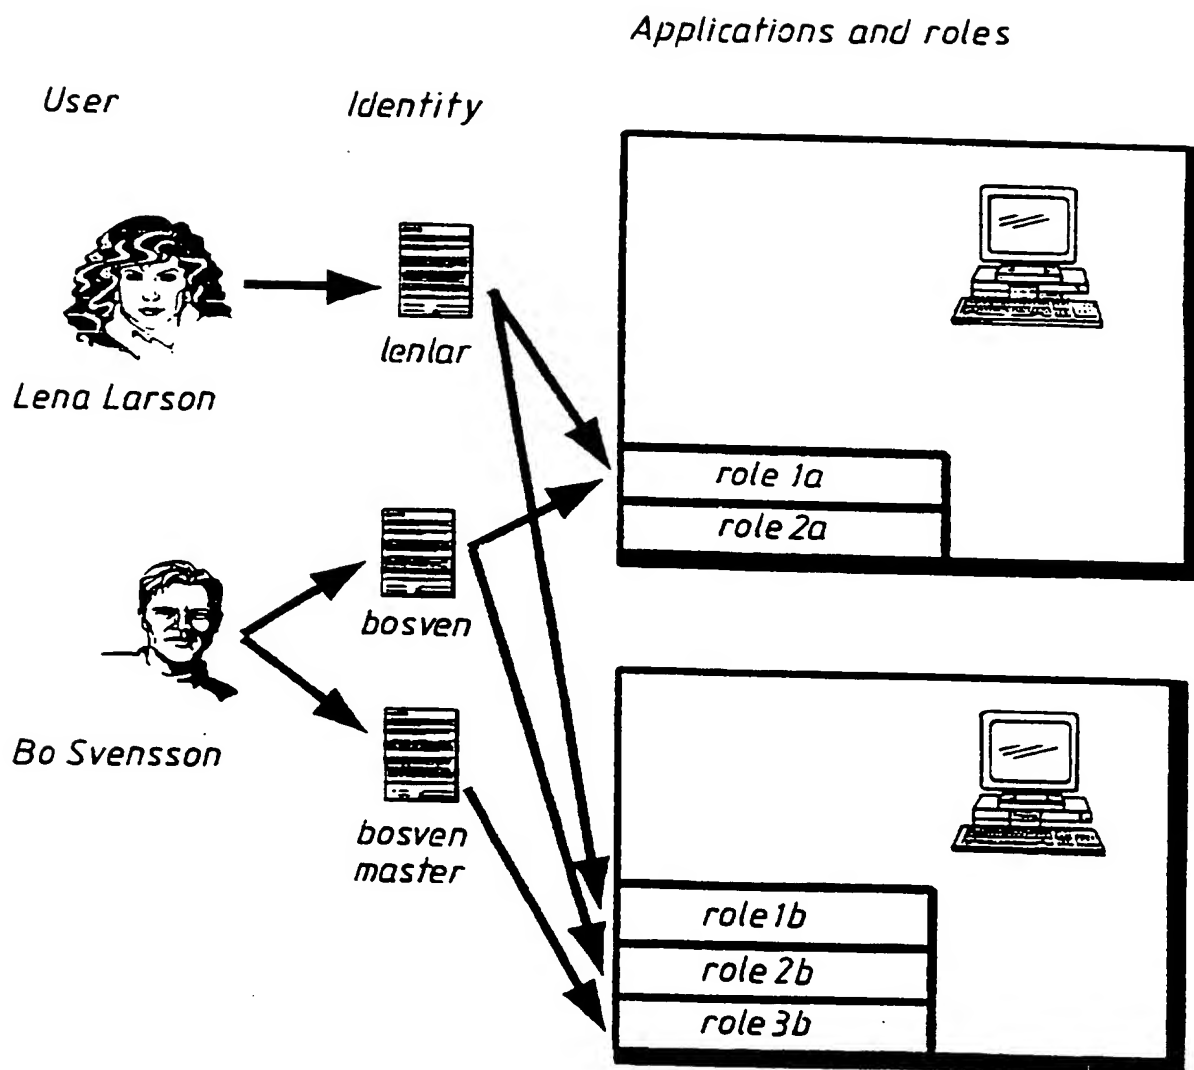
characterised by defining said roles in terms of system procedures which can be performed in relation to the different value sets held in said data base.

5 29. A method as claimed in claim 18, characterised by classifying said system procedures, for the purposes of role definition, by which of four basic data base operations, namely create, read, update and delete are applied by said system procedures to given value sets.

10 30. A method as claimed in any of claims 21 to 29, characterised by determining the access rights for a given user, in respect of a given data base application by running a data base procedure.

1 / 4

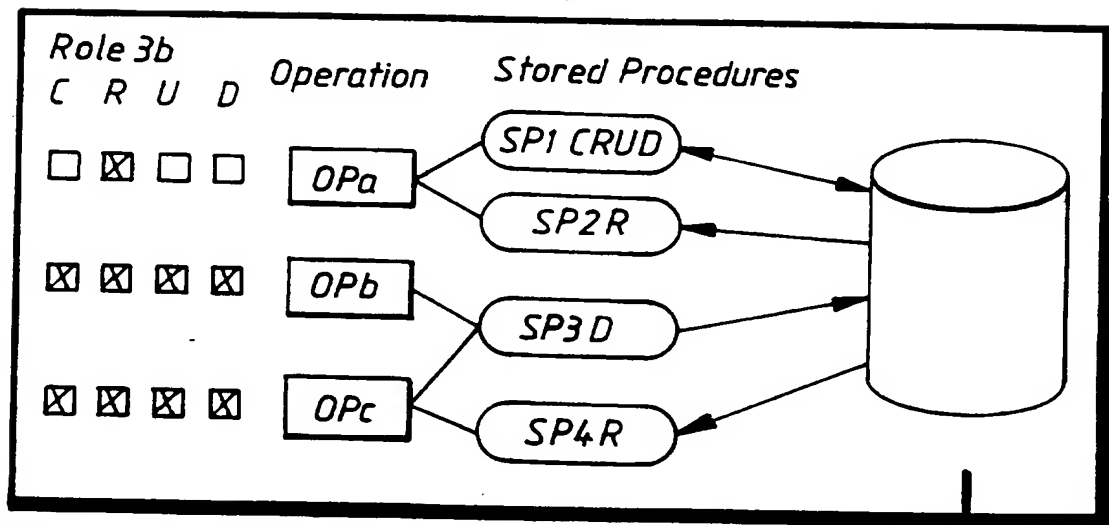
Fig. 1



2 / 4

Fig. 2

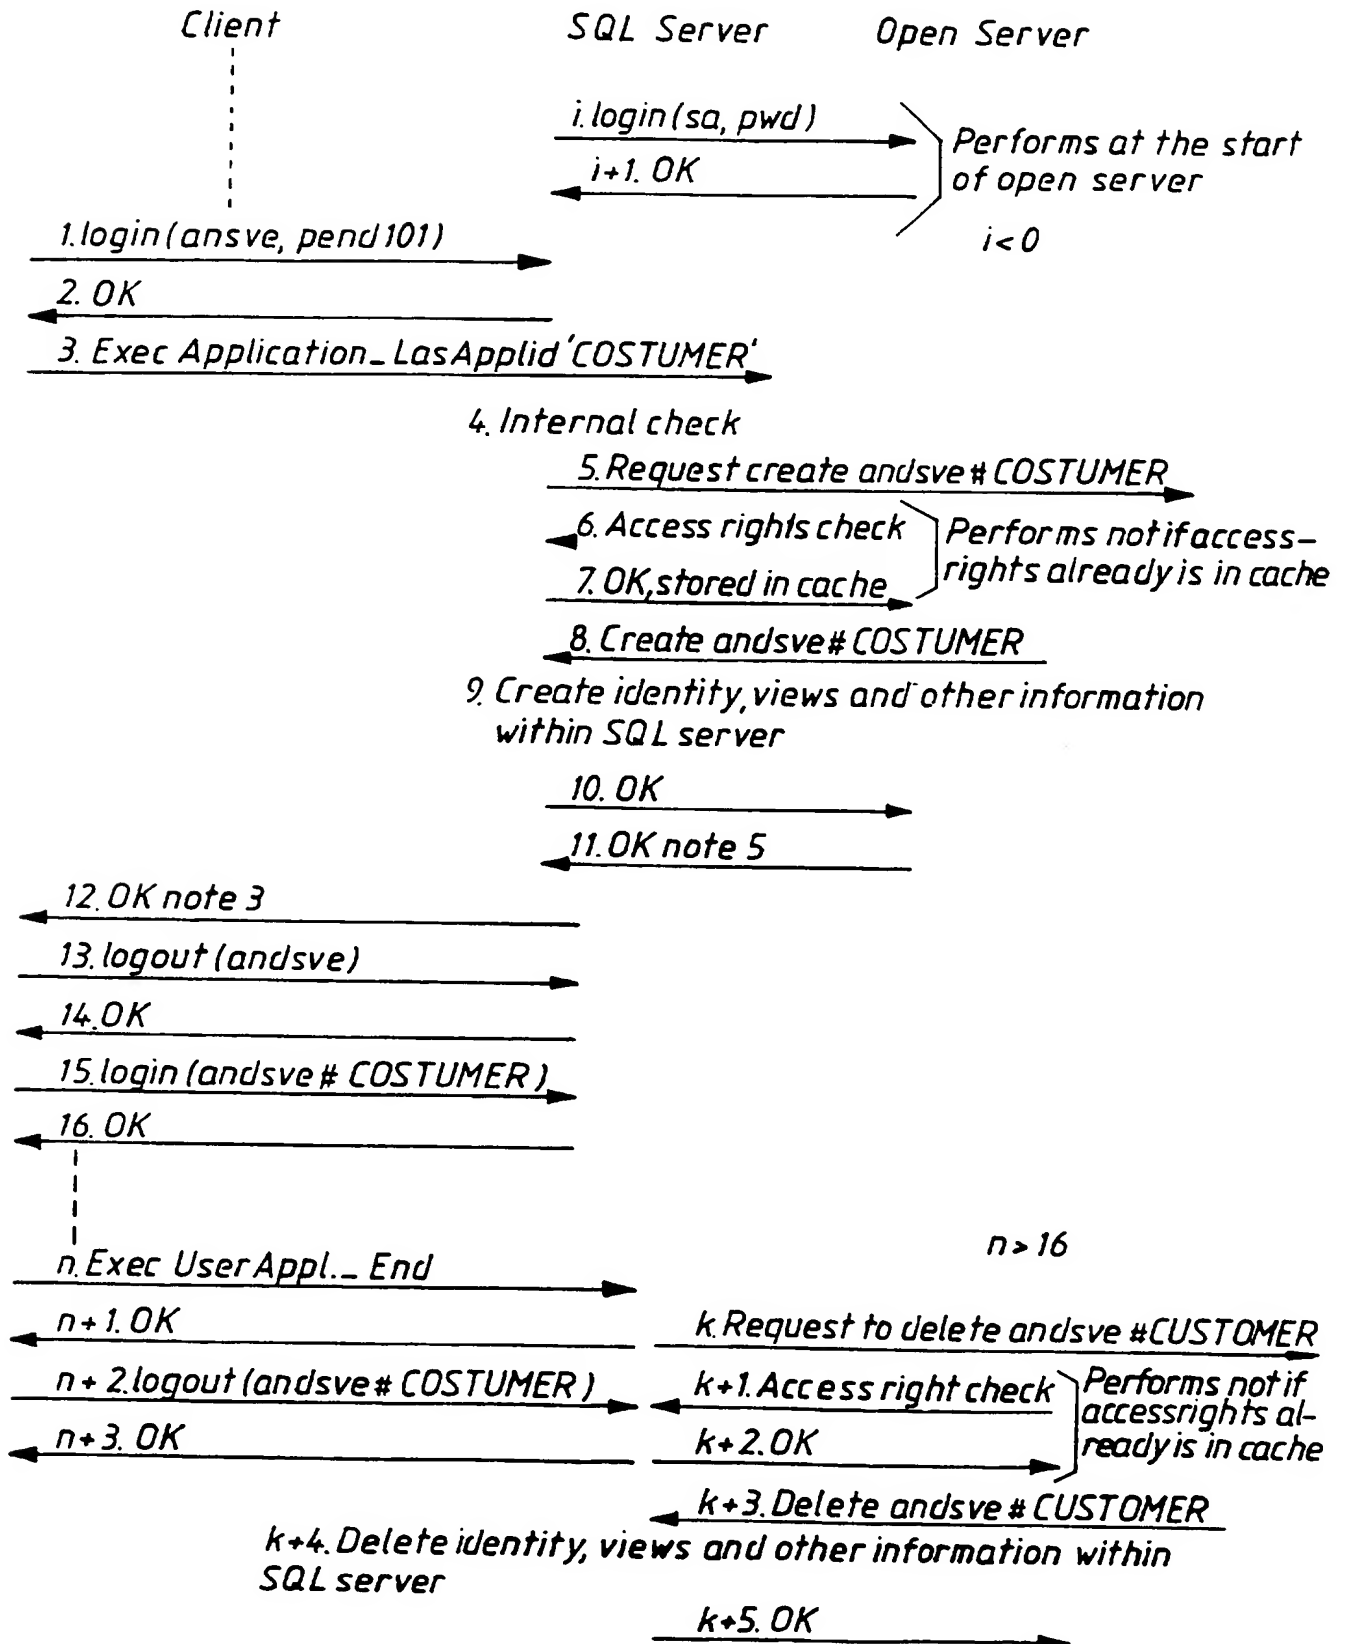
ROLE 3B

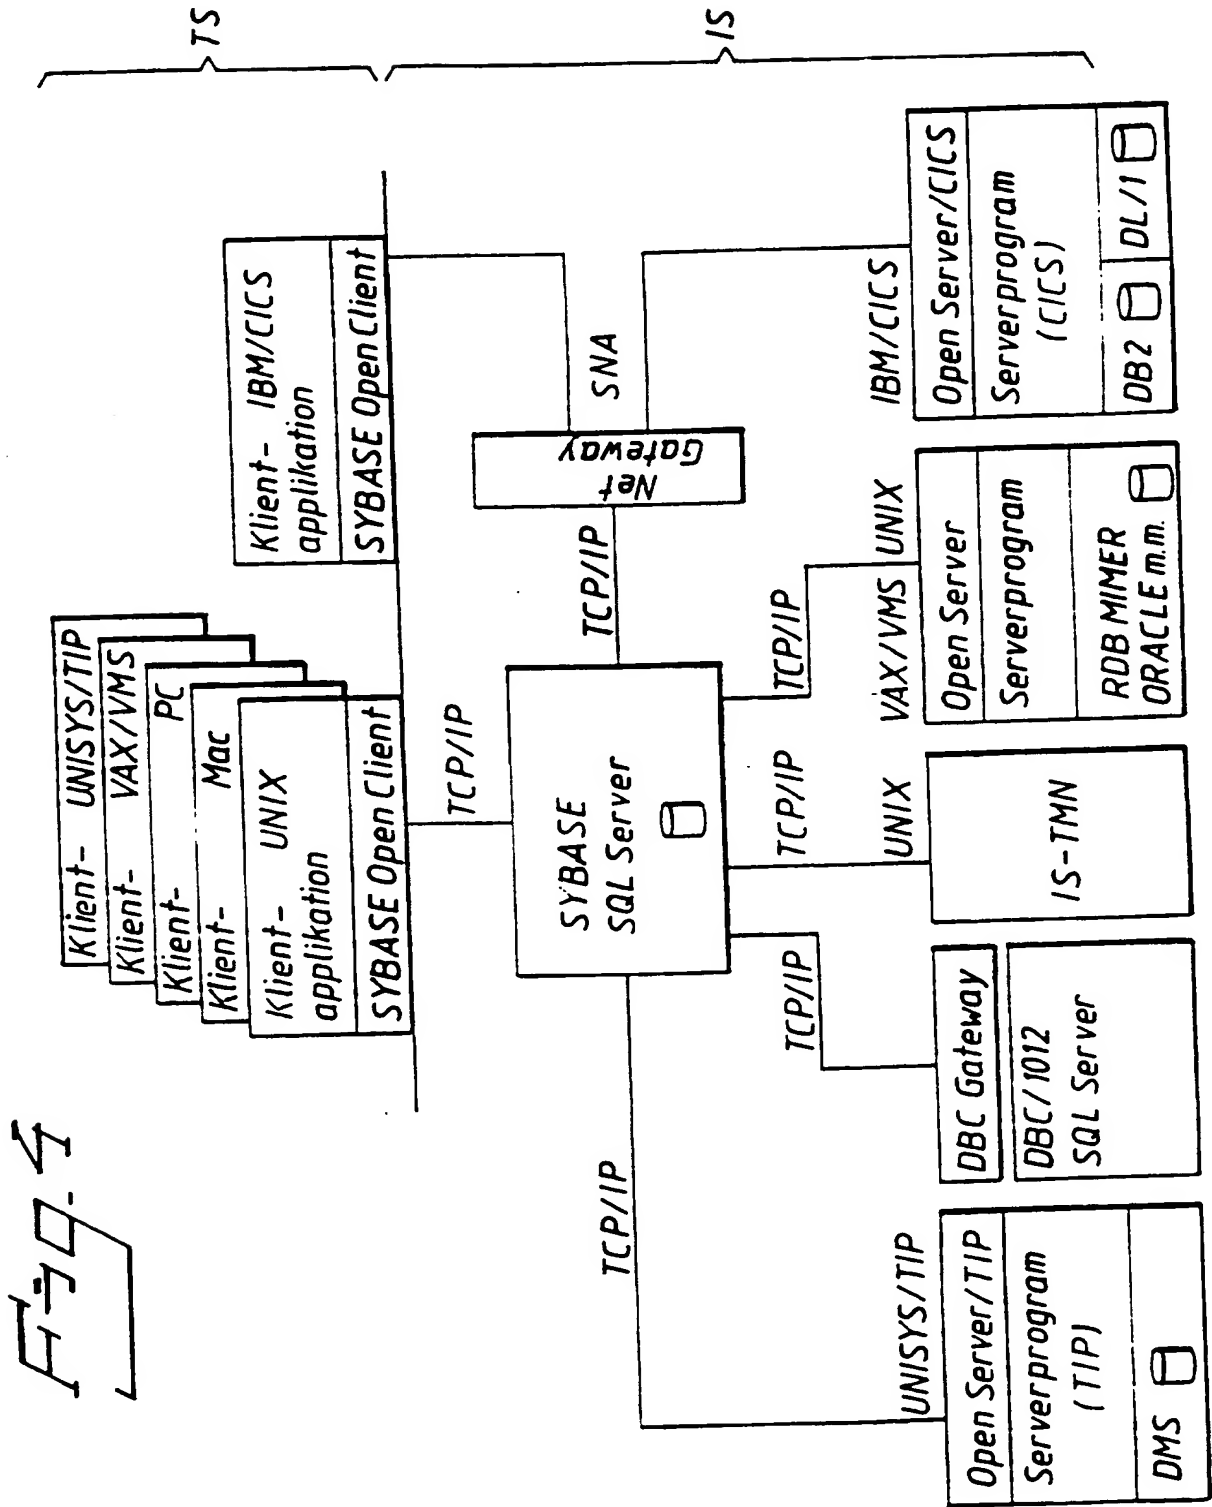


Role 3b	Value set
C R U D	
<input type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	Secret telephone nos
<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/>	Stockholm Region
<input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/> <input type="checkbox"/>	South Region
<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/>	Middle Region
<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/>	Northern Region

3 / 4

Fig. 3







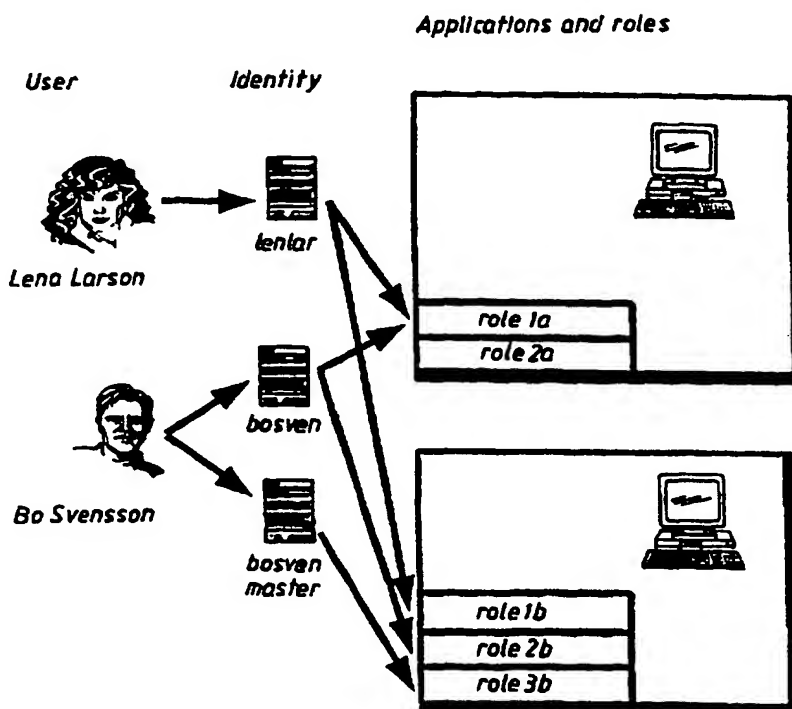
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 1/00, 12/14		A3	(11) International Publication Number: WO 96/17286
			(43) International Publication Date: 6 June 1996 (06.06.96)
(21) International Application Number: PCT/SE95/01394		(81) Designated States: US, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 22 November 1995 (22.11.95)		Published <i>With international search report.</i>	
(30) Priority Data: 9404157-1 29 November 1994 (29.11.94) SE		(88) Date of publication of the international search report: 29 August 1996 (29.08.96)	
(71) Applicant (for all designated States except US): TELIA AB [SE/SE]; S-123 86 Farsta (SE).			
(72) Inventor; and			
(75) Inventor/Applicant (for US only): BIGGE, Peter [SE/SE]; Styrmansgatan 17B, S-114 54 Stockholm (SE).			
(74) Agent: KARLSSON, Berne; Telia Research AB, Rudsjöterrassen 2, S-136 80 Haninge (SE).			

(54) Title: A METHOD FOR CONTROLLING ACCESS TO A DATA BASE, A DATA BASE AND A COMPUTER NETWORK USING THE SAME

(57) Abstract

The present invention solves the problem of access control in a large data base with a large number of users by defining access rights separately for each application run on the data base. A plurality of data base roles are defined, in terms of the basic data base procedures CRUD, (create, read, update and delete), which a user is entitled to use in respect of the various value sets stored in the data base. A given user is given access to a set of roles on an application-by-application basis. When a user logs-in to a particular data base application, to which he has access, a procedure is run by the file server on which the data base is stored. The file server generates the access rights for that user in respect of that application. The user rights are then stored in the file server, preferably on a cache memory, for the time in which the application in question is running on the file server. Preferably, when the user logs-out of that application the access rights are deleted from store.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Larvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

INTERNATIONAL SEARCH REPORT

1

International application No.

PCT/SE 95/01394

A. CLASSIFICATION OF SUBJECT MATTER

IPC6: G06F 1/00, G06F 12/14

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC6: G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

CLAIMS, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5347578 A (DUXBURY), 13 Sept 1994 (13.09.94), column 1, line 57 - column 5, line 40	1,2,4,5,11, 12,14,15,21, 22,24,25
A	--	6,16,26
Y	EP 0501475 A2 (BULL HN INFORMATION SYSTEMS INC.), 2 Sept 1992 (02.09.92), column 3, line 42 - line 56; column 4, line 40 - column 7, line 45	1,2,4,5,11, 12,14,15,21, 22,24,25
	--	

☒ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

7 June 1996

Date of mailing of the international search report

11-06- 1996

Name and mailing address of the ISA/
Swedish Patent Office
Box 5055, S-102 42 STOCKHOLM
Facsimile No. +46 8 666 02 86

Authorized officer

Bo Gustavsson

Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.

PCT/SE 95/01394

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	1990 IEEE Computer Society Symposium on Research in Security and Privacy, May 1990, Robert W. Baldwin, "Naming and Grouping Privileges to Simplify Security Management in large Databases", page 116 - page 132	1,2,4,5,11, 12,14,15,21, 22,24,25
A	--	6-10,16-20, 26-30
Y	1988 IEEE Symposium on Security and Privacy, April 1988, Stephen T. Vinter, "Extended Discretionary Access Controls", page 39 - page 49	1,2,4,5,11, 12,14,15,21, 22,24,25
A	-- -----	6-10,16-20, 26-30

INTERNATIONAL SEARCH REPORT

Information on patent family members

01/04/96

International application No.

PCT/SE 95/01394

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US-A- 5347578	13/09/94	AU-B- 658720 AU-A- 3527293 EP-A, A- 0561509 ZA-A- 9301487	27/04/95 23/09/93 22/09/93 04/10/93
EP-A2- 0501475	02/09/92	AU-B- 643366 AU-A- 1127092 JP-A- 5061833 US-A- 5274824	11/11/93 03/09/92 12/03/93 28/12/93

THIS PAGE BLANK (USPTO)